

IN THE SPECIFICATION:

On Page 1 line 5 - line 15, please replace the original paragraph with the following amended paragraph:

—

This application is related to the following United States Patent Applications:

Serial No. 09/933,883 filed on 8/20/2001 now issued as U. S. Patent No.
on entitled NEGOTIATED GRACEFUL TAKEOVER IN A NODE
CLUSTER, by Samuel M. Cramer et al.

Serial No. 09/933,866 filed on 8/20/2001 now issued as U. S. Patent No.
on entitled OPERATOR INITIATED GRACEFUL TAKEOVER IN A NODE
CLUSTER, by Samuel M. Cramer et al.

Serial No. 09/625,234 filed on 7/25/2000 now issued as U. S. Patent No.
6,728,897 on 4/27/2004 entitled NEGOTIATING TAKEOVER IN HIGH AVAILABILITY CLUSTER by Samuel M. Cramer, et al.

—

On Page 2 line 29 - Page 3 line 8 please replace the original paragraph with the following amended paragraph:

In one known file system, this ownership information is stored in two locations. Each disk has a predetermined sector arbitrarily named sector S that contains the definitive ownership information. In one example, sector S is sector zero of the disk. The second source of this ownership information is through the use of Small Computer System Interface (SCSI) level 3 reservations. These SCSI-3 reservations are described in SCSI Primary Commands – 3, by Committee T10 of the National Committee for Information Technology Standards, which is incorporated fully herein by reference. One technique for implementing disk ownership is described in commonly owned U.S. Patent Application, Serial No. 10/027,457, filed on 12/21/2001 now issued as U.S. Patent No. _____ on _____ entitled SYSTEM AND METHOD OF IMPLEMENTING DISK OWNERSHIP IN NETWORKED STORAGE by Susan M. Coatney, et al., which is hereby incorporated by reference.

--

On Page 3 lines 17 - 25 please replace the original paragraph with the following amended paragraph:

—

It is advantageous for the services and data provided by storage system to be available for access to the greatest degree possible. Accordingly, some computer storage systems provide a plurality of file servers (or filers) in a cluster, with a property that when a first filer fails, the second filer is available to take over and provide the services and the data otherwise provided by the first filer. When a first filer fails, the second filer in the cluster should assume the task of processing and handling any data access requests normally processed by the first filer. Such cluster configurations are described in United States Patent Application Serial No. 09/625,234 filed on 7/25/2000 now issued as U.S. Patent No. 6,728,897 on 4/27/2004 entitled NEGOTIATING TAKEOVER IN HIGH AVAILABILITY CLUSTER by Samuel M. Cramer, et al.

--

On Page 7 lines 24-26, please replace the original paragraph with the following amended paragraph:

--

Also attached to the network cloud 102 is a domain name system (DNS) server 130 132. The DNS server 132 provides name resolution services to other network devices in accordance with an appropriate name resolution protocol.

--

On Page 9 lines 19-25 please replace the original paragraph with the following amended paragraph:

--

The storage adapter 212 includes input/output (IO) interface circuitry that couples the filer to the disk shelves over an I/O interconnected arrangements such as a conventional high-speed high-performance fibre channel serial link topology ~~apology~~. The information is retrieved by the storage adapter, and if necessary, processed by the processor (or the adapter itself) prior to being forwarded over the system bus 214 to the appropriate network adapter 210 where the information is formatted into appropriate packets and returned to the client.

--

On Page 12 line 23 - Page 13 line 8, please replace the original paragraph with the following amended paragraph:

--

The resolvable name of the failover filer is resolved into the network address of the failover filer by the client. This resolution from a resolvable name to a network address is accomplished using known name resolution techniques. For example, in the network environment shown in Fig. 1, the resolvable name "FilerA-failover" would be resolved to the appropriate network address of Filer B, which is Filer A's clustered failover partner. Similarly, the name "FilerB-failover" would be resolved, using the appropriate name resolution service, to the network address of FilerA. Thus, a client can, by appending the set of string "-failover" to a known file server's name, ~~the client can~~ generate the resolvable name of the failover filer that is a cluster partner. From this computed resolvable name, the client can then generate the proper network address of the failover filer.

--

At Page 13 line 21 -Page 14 line 2, please replace the original paragraph with the following amended paragraph:

--

Fig. 4 shows the steps 400 of a procedure performed by the failover filer in accordance with this invention. In step 405, one of the filers in a clustered pair fails. This failure can be either due to a software and/or hardware problem, or due to a power failure and the like. In step 410, the failover filer detects the failure of its partner filer. This detection can occur by the failover filer detecting a lack of a heartbeat signal from the failed filer, or by other means, including an "I am failing" message transmitted by the failing filer over the cluster interconnect. The failover filer then takes over the failed filers disks, in step 415. This takeover of the failed filer's disks can be accomplished using known techniques to assert disk ownership in a network storage environment. One example is the method described in the above-incorporated patent application entitled SYSTEM AND METHOD OF IMPLEMENTING DISK OWNERSHIP IN NETWORKED STORAGE.

--

On Page 7 after line 3, please insert the following new paragraph:

--

Fig. 7 is a flow chart of a user initiated failover process.

--

At Page 17 lines 2-28, please replace the original paragraph with the following amended paragraph:

--

In certain instances, it is useful for a partner file server to take over the operations of another file server even though the other file server has not suffered an error condition. Such a user-initiated failover can be utilized when, for example, a repair or upgrade to the other file server needs to occur. These user initiated failovers are described generally in United States Patent Application Serial No. 09/933,883 filed on 8/20/2001 now issued as U.S. Patent No. _____ on _____ entitled NEGOTIATED GRACEFUL TAKEOVER IN A NODE CLUSTER, by Samuel M. Cramer et al. and in United States Patent Application Serial No. 09/933,866 filed on 8/20/2001 now issued as U.S. Patent No. _____ on _____ entitled OPERATOR INITIATED GRACEFUL TAKEOVER IN A NODE CLUSTER, by Samuel M. Cramer et al. By utilizing the principles of the present invention, no interruption in file services will occur to clients of the file server to be taken offline. The process 700 performed in such a user-initiated failover is shown in Fig. 7. In step 705, the user or administrator enters a failover command on the partner filer. Such a failover command would typically be included in the command set of the storage operating system associated with the file server. In response to the failover command,

the partner filer sends, in step 710, a "Please Shutdown" command to the other filer. Such a "Please Shutdown" command can be sent via the cluster interconnect. Next, in step 715, the partner filer starts a countdown timer. This countdown timer sets a time period for the partner filer to wait before attempting to take control and ownership of the other filers disks. The other filer precedes to shutdown (step 720) in response to the "Please Shutdown" command sent by the partner filer. At the expiration the countdown timer, the partner filer takes over the other filers disks in step 725. This taking over of ownership of disks is described in detail above. Next, in step 730, the partner filer activates its failover discriminator. After this failover discriminator has been activated, the partner filer can then service I/O requests from both its primary and secondary discriminators (step 740). The partner file server continues to service these I/O requests until such time that the user activates a giveback command in step 745. Upon initiation of the giveback command, the partner filer then performs a giveback operation (step 500).

--